

**RECEIVED
CENTRAL FAX CENTER****MAY 02 2007*****Gates & Cooper* LLP**Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045**FAX TRANSMISSION TO USPTO**TO: Commissioner for Patents
Attn: Examiner Cindy Nguyen
Patent Examining Corps
Facsimile Center
Alexandria, VA 22313-1450FROM: George H. Gates
OUR REF.: STL919990184US3
TELEPHONE: (310) 642-4146Total pages, including cover letter: 34PTO FAX NUMBER: 571-273-8300

If you do NOT receive all of the pages, please telephone us at (310) 641-8797, or fax us at (310) 641-8798.

Title of Document Transmitted:	TRANSMITTALS AND BRIEF OF APPELLANT
Applicant:	David E. Simmen
Serial No.:	10/807,871
Filed:	March 24, 2004
Group Art Unit:	2161
Title:	QUERY OPTIMIZATION TECHNIQUE FOR OBTAINING IMPROVED CARDINALITY ESTIMATES USING STATISTICS ON PRE-DEFINED QUERIES
Our Ref. No.:	STL919990184US3

Please charge all fees to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application.

By: Name: George H. Gates
Reg. No.: 33,500

I hereby certify that this paper is being transmitted by facsimile to the U.S. Patent and Trademark Office on the date shown below.


SignatureMay 2, 2007
Date

GHG/cf

G&C 30571.272-US-C1

Due Date: May 2, 2007

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**RECEIVED
CENTRAL FAX CENTER**

Applicant: David E. Simmen Examiner: Cindy Nguyen
Serial No.: 10/807,871 Group Art Unit: 2161
Filed: March 24, 2004 Docket: STL919990184US3
Title: QUERY OPTIMIZATION TECHNIQUE FOR OBTAINING IMPROVED CARDINALITY
ESTIMATES USING STATISTICS ON PRE-DEFINED QUERIES

MAY 02 2007

CERTIFICATE OF MAILING OR TRANSMISSION UNDER 37 CFR 1.8

I hereby certify that this correspondence is being filed via facsimile transmission to the U.S. Patent and Trademark Office on May 2, 2007.

By: George H. Gates
Name: George H. Gates

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

We are transmitting herewith the attached:

- ☒ Transmittal sheet, in duplicate, containing a Certificate of Mailing or Transmission under 37 CFR 1.8.
- ☒ Brief of Appellant(s).
- ☒ Charge the Fee for the Brief of Appellant(s) in the amount of \$500.00 to the Deposit Account.

Please consider this a **PETITION FOR EXTENSION OF TIME** for a sufficient number of months to enter these papers, if appropriate.

Please charge all fees to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application. A duplicate of this paper is enclosed.

Customer Number 45729
GATES & COOPER LLP
Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, CA 90045
(310) 641-8797

By: George H. Gates
Name: George H. Gates
Reg. No.: 33,500
GHG/cf

**RECEIVED
CENTRAL FAX CENTER**

MAY 02 2007

Due Date: May 2, 2007

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

In re Application of:)	
)	
Inventor: David E. Simmen)	Examiner: Cindy Nguyen
)	
Serial #: 10/807,871)	Group Art Unit: 2161
)	
Filed: March 24, 2004)	Appeal No.: _____
)	
Title: QUERY OPTIMIZATION TECHNIQUE)	
FOR OBTAINING IMPROVED)	
CARDINALITY ESTIMATES USING)	
STATISTICS ON PRE-DEFINED)	
QUERIES)	

BRIEF OF APPELLANT**MAIL STOP APPEAL BRIEF - PATENTS**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In accordance with 37 CFR §41.37, Appellant's attorney hereby submits the Brief of Appellant on appeal from the final rejection in the above-identified application, as set forth in the Office Action dated December 5, 2006.

Please charge the amount of \$500 to cover the required fee for filing this Appeal Brief as set forth under 37 CFR §41.37(a)(2) and 37 CFR §41.20(b)(2) to Deposit Account No. 09-0460 of IBM Corporation, the assignee of the present application.

Also, please charge any additional fees or credit any overpayments to Deposit Account No. 09-0460 of IBM Corporation.

I. REAL PARTY IN INTEREST

The real party in interest is International Business Machines Corporation, the assignee of the

05/03/2007 HLE333 00000092 090460 10807871

1

01 FC:1402 500.00 DA

present application.

II. RELATED APPEALS AND INTERFERENCES

There is a related appeal filed in U.S. Patent Application Serial No. 09/669,556, which is the parent of this application.

Appeals were filed in related U.S. Patent Nos. 6,738,755, 6,847,962 and 7,080,062, but a Board Decision was received only for U.S. Patent No. 7,080,062 (the other patents were allowed after the filing of an Appeal Brief). A copy of that Board Decision is attached in the Related Proceedings Appendix.

III. STATUS OF CLAIMS

Claims 1-33 are pending in the application.

Claims 1-33 were provisionally rejected on the grounds of nonstatutory double patenting over claim 1 and 3-33 of copending Application No. 09/669,556.

Claims 1, 2-5, 12-17, and 23-28 were rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Publication No. 2003/0088558 (Zaharioudakis).

Claims 6-11, 17-22, and 28-33 were rejected under 35 U.S.C. §103(a) as being unpatentable over Zaharioudakis in view of U.S. Patent No. 6,496,819 (Bello).

Claims 1-33 are being appealed.

IV. STATUS OF AMENDMENTS

A response under 37 C.F.R. §1.116 was submitted subsequent to the final Office Action, on February 1, 2007, but no claims were amended in the response. An Advisory Action was mailed on February 21, 2007 stating that the response overcame the rejections, but then stated a different grounds for rejection based on Provisional Application Serial No. 60/135,133.

A response under 37 C.F.R. §41.33(a) is being submitted subsequent to the final Office Action, on the same date as this Brief. The response under 37 C.F.R. §41.33(a) includes a Terminal Disclaimer to overcome the obviousness-type double patenting rejections of claims 1 and 3-33.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Appellant's invention, as recited in independent claims 1, 12 and 23, are generally directed to optimizing execution of a query that accesses data stored on a data store connected to a computer.

Independent claim 1 recites a method of optimizing execution of a query that accesses data stored on a data store (104) connected to a computer (102). (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1.) The method includes the step of using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query. (See page 3, lines 19-29; page 6, lines 18-22 referring to 124 in FIG. 1; page 14, line 19 – page 18, line 7, referring to 124 in FIG. 1; page 18, line 9 – page 29, line 30, referring to 124 in FIG. 1; and page 30, line 2 - page 31, line 20 referring 200-204, 300-310, and 400-416 in FIGS. 2, 3, 4A and 4B, respectively.) In addition, the method includes the step of executing the optimal query execution plan for the query in order to access the data stored on the data store (104) connected to a computer (104) and then output the accessed data. (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1.)

Independent claim 12 recites an apparatus for optimizing execution of a query. The apparatus includes a computer (102) having a data store (104) coupled thereto, wherein the data store (104) stores data. (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1.) In addition, the apparatus includes one or more computer programs (108-124), performed by the computer (102), for using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query, and executing the optimal query execution plan for the query in order to access the data stored on the data store (104) connected to a computer (102) and then output the accessed data. (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1; page 6, lines 18-22 referring to 124 in FIG. 1; page 14, line 19 – page 18, line 7, referring to 124 in FIG. 1; page 18, line 9 – page 29, line 30, referring to 124 in FIG. 1; and page 30, line 2 - page 31, line 20 referring 200-204, 300-310, and 400-416 in FIGS. 2, 3, 4A and 4B, respectively.)

Independent claim 23 recites an article of manufacture comprising a program storage medium readable by a computer (102) and embodying one or more instructions executable by the

computer (102) to optimizing execution of a query that accesses data stored on a data store (104) connected to the computer (102). (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1.) The instructions include using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query. (See page 3, lines 19-29; page 6, lines 18-22 referring to 124 in FIG. 1; page 14, line 19 – page 18, line 7, referring to 124 in FIG. 1; page 18, line 9 – page 29, line 30, referring to 124 in FIG. 1; and page 30, line 2 – page 31, line 20 referring 200-204, 300-310, and 400-416 in FIGS. 2, 3, 4A and 4B, respectively.) In addition, the instructions include executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data. (See page 3, lines 19-29; page 4, lines 25-27 referring to 102 in FIG. 1; page 5, lines 13-28 referring to 102-108 and 114-124 in FIG. 1.)

VI. GROUND OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether claims 1, 2-5, 12-17, and 23-28 are anticipated under 35 U.S.C. §102(e) by U.S. Publication No. 2003/0088558 (Zaharioudakis).

2. Whether claims 6-11, 17-22, and 28-33 are obvious under 35 U.S.C. §103(a) over U.S. Publication No. 2003/0088558 (Zaharioudakis) in view of U.S. Patent No. 6,496,819 (Bello).

VII. ARGUMENTS

A. Arguments directed to the first grounds for rejection: Whether claims 1, 2-5, 12-17, and 23-28 are anticipated under 35 U.S.C. §102(e) by U.S. Publication No. 2003/0088558 (Zaharioudakis).

1. Claims 1, 12 and 23

On page (2) of the Office Action, claims 1, 2-5, 12-17, and 23-28 were rejected under 35 U.S.C. §102(e) as being anticipated by U.S. Publication No. 2003/0088558 (Zaharioudakis).

Appellant's attorney respectfully traverses these rejections. Specifically, Appellant's attorney asserts that Zaharioudakis is not a prior art reference for the following reasons:

1. This application is a continuation of Utility Application Serial No. 09/669,556, filed on September 26, 2000, which application claims the benefit of Provisional Application Serial No. 60/171,797, filed on December 22, 1999.

2. Zaharioudakis has a November 5, 2002 filing date, which is more than two years after the September 26, 2000 filing date of the parent Utility Application Serial No. 09/669,556, and which is more than three years after the December 22, 1999 filing date of the parent Provisional Application Serial No. 60/171,797.
3. Although Zaharioudakis is a continuation-in-part to Utility Application Serial No. 09/502,821, filed on February 11, 2000, which claims priority to Provisional Application Serial No. 60/135,133, filed on May 20, 1999, at least some of the portions of Zaharioudakis referred to in the rejections of the Office Action are only entitled to the November 5, 2002 filing date, because these portions of Zaharioudakis cannot be found in any of the prior patent applications. Note, for example, that the rejections of independent claims 1, 12 and 23 refer to the following portions of Zaharioudakis: paragraphs [0041], [0042] and [0043]. However, Appellant's attorney submits that paragraphs [0042]-[0043] of Zaharioudakis cannot be found in any of the prior patent applications, paragraphs [0042]-[0043] of Zaharioudakis are not entitled to an effective date earlier than the November 5, 2002 filing date, and thus paragraphs [0042]-[0043] of Zaharioudakis cannot be cited against Appellant's claims.

Consequently, the Office Action fails to comply with the requirements of M.P.E.P. §706.02(f)(1), which states that "[t]he 35 U.S.C. 102(e) date of a reference ... is its earliest effective U.S. filing date, taking into consideration any proper benefit claims to prior U.S. applications under 35 U.S.C. 119(e) or 120 if the prior application(s) properly supports the subject matter used to make the rejection in compliance with 35 U.S.C. 112, first paragraph."

After making these arguments in the response to the final rejection set forth in the Office Action dated December 5, 2006, the Examiner issued an Advisory Action stating that Appellant's reply had overcome the rejections. Nonetheless, in the Advisory Action, the Examiner stated the following:

Applicant's reply has overcome the following rejection(s): Applicant argue that Zaharioudakis is not a prior art preference because Zaharioudakis has a November 5, 2002 filing date, which is more than two years after the September 26, 2000 filing date of the present application, and which is more than three years after the December 22, 1999 priority date of the present application. Although Zaharioudakis is a continuation-in-part to Utility Application Serial No. 09/502,821, filed on February 11, 2000, which claims priority to Provisional Application Serial

No. 60/135,133, filed on May 20, 1999, at least some of the specific portions of Zaharioudakis referred to in the rejections of the Office Action are only entitled to the November 5, 2002 filing date, because these specific portions of Zaharioudakis cannot be found in any of the prior patents. Note, for example, that the rejections of independent claims 1, 11 and 21 refer to the following portions of Zaharioudakis: paragraphs [0031], [0041] and [0043]. However, Applicant's attorney submits that paragraph [0043] of Zaharioudakis cannot be found in any of the prior patents, paragraph [0043] of Zaharioudakis is not entitled to an effective date earlier than the November 5, 2002 filing date, and thus paragraph [0043] of Zaharioudakis cannot be cited against Applicant's claims.

In response, the provisional Application serial number 60/135133, filed on May 20, 1999 provided all the portions of Zaharioudakis as cited in the rejection as following:

Regarding claims 1, 12 and 23, Zaharioudakis discloses: a method, an apparatus, an article of manufacture for optimizing execution of a query that accesses data stored on a data store connected to a computer comprising: using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query (system held statistics on the data to be access as the size of the table, the number of distinct values in particular column, page 9, lines 20-24 and matching between query and AST, the query can be optimized by re-writing it to use the AST, page 15, lines 14 to page 16, lines 17) or (0041-0043, Zaharioudakis); executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data. (page 16, lines 15-18, Provisional application) or (paragraphs 0041, 0043, Zaharioudakis).

Appellant's attorney interprets the Advisory Action as maintaining the rejection based on paragraphs [0041]-[0043] of Zaharioudakis (notwithstanding the initial assertion to the contrary), and as improperly introducing a new grounds for rejection based on Provisional Application Serial No. 60/135,133.

First, Appellant's attorney respectfully disagrees that paragraphs [0041]-[0043] of Zaharioudakis are provided by the cited portions of Provisional Application Serial No. 60/135,133.

Consider paragraphs [0041]-[0043] of Zaharioudakis, which are set forth in their entirety below:

[0041] FIG. 3 is a flowchart illustrating the steps necessary for the interpretation and execution of SQL statements embedded in source code according to an embodiment of the present invention. Block 300 represents program source code containing a host language (such as COBOL or C) and embedded SQL statements. The program source code is then input to a pre-compile step 302. There are two outputs from the pre-compile step 302: a modified source module 304 and a

Database Request Module (DBRM) 306. The modified source module 304 contains host language calls to the RDBMS software, which the pre-compile step 302 inserts in place of SQL statements. The DBRM 306 is comprised of the SQL statements from the program source code 300. A compile and link-edit step 308 uses the modified source module 304 to produce a load module 310, while an optimize and bind step 312 uses the DBRM 306 to produce a compiled set of runtime structures for the application plan 314. As indicated above in conjunction with FIG. 2, the SQL statements from the program source code 300 specify only the desired data, but not how to retrieve the data. The optimize and bind step 312 may optimize the SQL query in a manner described in more detail later in this specification. Thereafter, the optimize and bind step 312 considers both the available access paths (indexes, sequential reads, etc.) and system held statistics on the data to be accessed (the size of the table, the number of distinct values in a particular column, etc.), to choose what it considers to be the most efficient access path for the query. The load module 310 and application plan 314 are then executed together at step 316.

[0042] 4 An Optimization Technique for Automatic Summary Tables

[0043] Automatic summary tables (ASTs) (also known as materialized views) are database tables that contain the pre-computed results of certain SQL queries. The purpose of the optimization technique presented here is to take advantage of existing ASTs by employing a matching algorithm that determines whether the content of an AST overlaps with the content of an SQL query, and compensates for the non overlapping parts. When such an overlap exists, we say that the query and the AST match. After discovering a match, the optimizer may choose to rewrite the query so that it will access the AST instead of one or more of the base tables. The optimizer will rewrite the query if the estimated cost of the rewritten query is less than the original query.

On the other hand, page 9, lines 20-24 and page 15, lines 14 to page 16, lines 17 of Provisional Application Serial No. 60/135,133 recite the following:

Page 9, lines 20-24 (actually lines 19-23)

Thereafter, the optimize and bind step 314 considers both the available access paths (indexes, sequential reads, etc.) and system held statistics on the data to be accessed (the size of the table, the number of distinct values in a particular column, etc.), to choose what it considers to be the most efficient access path for the query. The load module 312 and application plan 316 are then executed together at step 318.

Page 15, line 14 through page 16, line 17

4. Matching Between Query and AST

It can be said that a query "matches" with an AST if there is some box of the query QGM graph that "matches" with the top (root) box of the AST QGM graph.

If this is true, then the query can be optimized by re-writing it to use the AST. This is shown in FIG. 6, which is a flowchart illustrating the method of optimizing SQL queries in step 204 of FIG. 2 and step 314 of FIG. 3 according to the preferred embodiment of the present invention.

Block 600 represents the computer system 100, specifically an optimizer function of the RDBMS software 106, accepting a query.

Block 602 is a decision block that represents the computer system 100 determining whether there is one or more summary tables referencing one or more tables in the query. If so, control transfers to Block 604; otherwise, control transfers to Block 610.

Block 604 represents the computer system 100 analyzing whether a summary table can be used to answer the query. Specifically, this Block performs subsumption tests between the query and the definition of the summary table.

Block 606 is a decision block that represents the computer system 100 determining whether the query should be rewritten to take advantage of one or more of the summary tables. If so, control transfers to Block 608; otherwise, control transfers to Block 610.

Block 608 represents the computer system 100 rewriting the query to use the identified summary tables for answering the query. Specifically, this Block compensates complex expressions using the summary table as they are identified in Block 604, wherein the expressions can be re-derived from one or more of the columns of the summary table.

Block 610 represents the computer system 100 executing the query.

After these query transformation steps are performed, block 612 returns control to block 204 in FIG. 2 or block 314 in FIG. 3 for subsequent processing steps, including the execution of the SQL query against the relational database and the output of the result set.

There is a lack of identity between paragraphs [0042]-[0043] of Zaharioudakis and the cited portions of Provisional Application Serial No. 60/135,133. In view of this fact, Appellant's attorney again submits that paragraphs [0042]-[0043] of Zaharioudakis cannot be found in Provisional Application Serial No. 60/135,133, paragraphs [0042]-[0043] of Zaharioudakis are not entitled to the filing date of Provisional Application Serial No. 60/135,133 or any date earlier than the November 5, 2002 filing date of Zaharioudakis, and thus paragraphs [0042]-[0043] of Zaharioudakis cannot be cited against Appellant's claims.

Moreover, Appellant's attorney submits that paragraphs [0041]-[0043] of Zaharioudakis do not teach or suggest Appellant's claims, namely optimizing execution of a query by "using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query."

Instead, paragraphs [0041]-[0043] of Zaharioudakis merely describe the use of ASTs (automatic summary tables). However, there is no discussion in paragraphs [0041]-[0043] of Zaharioudakis of using statistics on expressions of pre-defined queries to determine an optimal query execution plan. Indeed, nowhere do paragraphs [0041]-[0043] of Zaharioudakis describe maintaining such statistics or using such statistics to optimize queries. Instead, the only statistics are described in paragraph [0041] of Zaharioudakis and merely comprise the size of the table and the number of distinct values in a particular column.

In addition, although Appellant's attorney submits that the cited portions of Provisional Application Serial No. 60/135,133 are improperly introduced for the first time in the Advisory Action, rather than an Office Action, Appellant's attorney will also address the substance of the cited portions of Provisional Application Serial No. 60/135,133.

Specifically, Appellant's attorney submits that the cited portions of Provisional Application Serial No. 60/135,133 do not teach or suggest Appellant's claims, namely optimizing execution of a query by "using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query."

Instead, the cited portions of Provisional Application Serial No. 60/135,133 merely describe the use of ASTs. However, there is no discussion in the cited portions of Provisional Application Serial No. 60/135,133 of using statistics on expressions of pre-defined queries to determine an optimal query execution plan. Indeed, nowhere do the cited portions of Provisional Application Serial No. 60/135,133 describe maintaining such statistics or using such statistics to optimize queries. Instead, the only statistics described in the cited portions of Provisional Application Serial No. 60/135,133 are the size of the table and the number of distinct values in a particular column.

Consequently, Appellants' attorney submits that independent claims 1, 12 and 23 are allowable over Zaharioudakis or Provisional Application Serial No. 60/135,133. Further, dependent claims 2-5, 13-17, and 24-28 are submitted to be allowable over Zaharioudakis or Provisional Application Serial No. 60/135,133 in the same manner, because they are dependent on independent claims 1, 12 and 23, respectively, and thus contain all the limitations of independent claims 1, 12 and 23. In addition, dependent claims 2-5, 13-17, and 24-28 recite limitations not shown by Zaharioudakis or Provisional Application Serial No. 60/135,133, as described in more detail below.

2. Claims 2, 13 and 24

With regard to dependent claims 2, 13, and 24, which recite that each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view, these claims stand or fall with independent claims 1, 12 and 23.

3. Claims 3, 14 and 25

With regard to dependent claims 3, 14, and 25, which recite generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query, and using the generated cardinality estimates to determine an optimal query execution plan for the query, the Office Action asserts that these limitations are described in paragraph [0043] of Zaharioudakis. Appellant's attorney disagrees. Paragraph [0043] of Zaharioudakis is not prior art to these claims. Moreover, paragraph [0043] of Zaharioudakis says nothing about generating cardinality estimates for query execution plans using the statistics of the pre-defined queries that vertically overlap the query, or using the generated cardinality estimates to determine an optimal query execution plan for the query. Similarly, the cited portions of Provisional Application Serial No. 60/135,133 do not teach or suggest these limitations.

4. Claims 4, 15 and 26

With regard to dependent claims 4, 15, and 26, which recite that the statistics are used to improve a combined selectivity estimate of one or more predicates of the query, the Office Action asserts that these limitations are described in paragraph [0041] of Zaharioudakis. Appellant's attorney disagrees. Paragraph [0041] of Zaharioudakis merely describes statistics comprising the size of the table and the number of distinct values in a particular column. However, paragraph [0041] of Zaharioudakis does not teach or suggest using statistics on expressions of pre-defined queries to determine an optimal query execution plan, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query. Similarly, the cited portions of Provisional Application Serial No. 60/135,133 do not teach or suggest these limitations.

5. Claims 5, 16 and 27

With regard to dependent claims 5, 16, and 27, which recite that predicates are applied by one or more of the pre-defined queries, these claims stand or fall with independent claims 1, 12 and 23.

B. Arguments directed to the second grounds for rejection: Whether claims 6-11, 17-22, and 28-33 are obvious under 35 U.S.C. §103(a) over U.S. Publication No. 2003/0088558 (Zaharioudakis) in view of U.S. Patent No. 6,496,819 (Bello).

On page (4) of the Office Action, claims 6-11, 17-22, and 28-33 were rejected under 35 U.S.C. §103(a) as being unpatentable over Zaharioudakis in view of U.S. Patent No. 6,496,819 (Bello).

Appellant's attorney respectfully traverses these rejections. Specifically, Appellant's attorney asserts that Zaharioudakis is not a prior art reference for the reasons set forth above. Moreover, Appellant's attorney submits that the cited portions of Provisional Application Serial No. 60/135,133 are improperly introduced for the first time in the Advisory Action, rather than an Office Action.

Notwithstanding these facts, Appellant's attorney submits that the cited portions of Zaharioudakis or Provisional Application Serial No. 60/135,133, when combined with U.S. Patent No. 6,496,819 (Bello), do not teach or suggest Appellant's dependent claims 6-11, 17-22, and 28-33.

1. Claims 6, 17 and 28

With regard to dependent claims 6, 17 and 28, which recite that the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query, the Office Action admits that Zaharioudakis does not disclose these limitations, but asserts that Bello discloses these limitations at col. 10, lines 45-76, and that it would have been obvious to combine the references to enable the query reduction factor to estimates how useful it will be to access the materialized view to process the received query.

Appellant's attorney disagrees. The cited portions of Bello merely disclose a query reduction factor, which is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view. However, the query reduction factor is not the same as the selectivity estimate, which comprises a ratio of the

cardinality of the pre-defined query to the product of cardinalities of base tables referenced in the pre-defined query, which is a different value. Moreover, the motivation to combine the references is suggested by the Office Action, not by either of the references, and therefore constitutes improper hindsight.

2. Claims 7, 18 and 29

With regard to dependent claims 7, 18 and 29, which recite that zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query, the Office Action admits that Zaharioudakis does not disclose these limitations, but asserts that Bello discloses these limitations at col. 10, lines 20-45, and that it would have been obvious to combine the references to enable the system determines whether the materialized view is actually eligible to be used in rewrite of the received query to reduce the execution cost of the query.

Appellant's attorney disagrees. The cited portions of Bello merely disclose that the database server determines whether a materialized view is eligible to be used in a rewrite of a query, but not whether zero or more predicates of the query are applied by a pre-defined query and the remaining predicates of the query are eligible to be applied on the pre-defined query. Moreover, the motivation to combine the references is suggested by the Office Action, not by either of the references, and therefore constitutes improper hindsight.

3. Claims 8, 19 and 30

With regard to dependent claims 8, 19 and 30, which recite that a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query, these claims stand or fall with claims 7, 18 and 29.

4. Claims 9, 20 and 31

With regard to dependent claims 9, 20 and 31, which recite determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query that zero or more predicates of the query are applied by one of the pre-defined

queries and wherein the remaining predicates are eligible to be applied on the pre-defined query, the Office Action asserts that Bello discloses these limitations at col. 10, lines 30-36.

Appellant's attorney disagrees. The cited portions of Bello merely disclose that an unprocessed possible materialized view is selected, and a determination is made whether the selected materialized view is eligible for use in rewriting the received query. However, nothing in the cited portions of Bell teach or suggest determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query that zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

5. Claims 10, 21 and 32

With regard to dependent claims 10, 21 and 32, which recite that a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query, the Office Action asserts that Bello discloses these limitations at col. 10, lines 37-56.

Appellant's attorney disagrees. The cited portions of Bello merely disclose that a query reduction factor, which is the ratio of (1) the sum of the cardinalities of matching relations in the query that will be replaced by the materialized view to (2) the cardinality of the materialized view. However, the query reduction factor is not the same as the cardinality ratio, which comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query. Instead, these are completely different values.

6. Claims 11, 22 and 33

With regard to dependent claims 11, 22 and 33, which recite that the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio, the Office Action asserts that Bello discloses these limitations at col. 11, lines 55 to col. 12, lines 41.

Appellant's attorney disagrees. The cited portions of Bello merely describe a join compatibility test, which is an eligibility test performed by a database server to determine whether a materialized view can be used to process a received query. However, the join compatibility test of

Bello is in no way related to a selectivity estimate that comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

VIII. CONCLUSION

In light of the above arguments, Appellant's attorney respectfully submits that the cited references do not anticipate nor render obvious the claimed invention. More specifically, Appellant's claims recite novel physical features which patentably distinguish over any and all references under 35 U.S.C. §§ 102 and 103.

As a result, a decision by the Board of Patent Appeals and Interferences reversing the Examiner and directing allowance of the pending claims in the subject application is respectfully solicited.


Respectfully submitted,

GATES & COOPER LLP
Attorneys for Appellant

Howard Hughes Center
6701 Center Drive West, Suite 1050
Los Angeles, California 90045
(310) 641-8797

Date: May 2, 2007

GHG/

By: 
Name: George H. Gates
Reg. No.: 33,500

CLAIMS APPENDIX

1. A method of optimizing execution of a query that accesses data stored on a data store connected to a computer, comprising:
using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query; and
executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.
2. The method of claim 1, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.
3. The method of claim 1, further comprising:
generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and
using the generated cardinality estimates to determine an optimal query execution plan for the query.
4. The method of claim 3, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.
5. The method of claim 4, wherein the predicates are applied by one or more of the pre-defined queries.
6. The method of claim 5, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.
7. The method of claim 4, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

8. The method of claim 7, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

9. The method of claim 8, further comprising determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

10. The method of claim 9, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

11. The method of claim 10, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

12. An apparatus for optimizing execution of a query, comprising:
a computer having a data store coupled thereto, wherein the data store stores data;
one or more computer programs, performed by the computer, for using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query, and executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.

13. The apparatus of claim 12, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.

14. The apparatus of claim 12, further comprising:
one or more computer programs for generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and
one or more computer programs for using the generated cardinality estimates to determine an optimal query execution plan for the query.

15. The apparatus of claim 14, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.

16. The apparatus of claim 15, wherein the predicates are applied by one or more of the pre-defined queries.

17. The apparatus of claim 16, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

18. The apparatus of claim 15, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

19. The apparatus of claim 18, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

20. The apparatus of claim 19, further comprising one or more computer programs for determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

21. The apparatus of claim 20, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

22. The apparatus of claim 21, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

23. An article of manufacture comprising a program storage medium readable by a computer and embodying one or more instructions executable by the computer to optimizing execution of a query that accesses data stored on a data store connected to the computer, comprising:

using statistics on one or more expressions of one or more pre-defined queries to determine an optimal query execution plan for the query; and

executing the optimal query execution plan for the query in order to access the data stored on the data store connected to a computer and then output the accessed data.

24. The article of claim 23, wherein each of the pre-defined queries is associated with an automatic summary table, a materialized view or a view.

25. The article of claim 23, further comprising:
generating cardinality estimates for one or more query execution plans for the query using the statistics of one or more of the pre-defined queries that vertically overlap the query; and
using the generated cardinality estimates to determine an optimal query execution plan for the query.

26. The article of claim 25, wherein the statistics are used to improve a combined selectivity estimate of one or more predicates of the query.

27. The article of claim 26, wherein the predicates are applied by one or more of the pre-defined queries.

28. The article of claim 27, wherein the selectivity estimate comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

29. The article of claim 26, wherein zero or more predicates of the query are applied by one of the pre-defined queries and wherein the remaining predicates are eligible to be applied on the pre-defined query.

30. The article of claim 29, wherein a predicate is eligible to be applied on the pre-defined query if it can be evaluated using the output columns and expressions of the pre-defined query.

31. The article of claim 30, further comprising determining a subpredicate combined selectivity estimate of the unapplied eligible predicates using column distribution statistics of the pre-defined query.

32. The article of claim 31, wherein a cardinality ratio comprises a ratio of a cardinality of the pre-defined query to a product of cardinalities of base tables referenced in the pre-defined query and the query.

33. The article of claim 32, wherein the selectivity estimate comprises a product of the subpredicate combined selectivity estimate and the cardinality ratio.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

Enclosed herewith is a copy of the "Decision on Appeal" received from the U.S. Patent and Trademark Office (PTO) in connection with U.S. Patent Application Serial No. 09/502,820, now U.S. Patent No. 7,080,062.

The opinion in support of the decision being entered today was not written for publication and is not binding precedent of the Board.

Paper No. 30


UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES

Ex parte TING YU LEUNG,
DAVID E. SIMMEN, and YANG SUN

Appeal No. 2004-1538
Application 09/502,820¹

ON BRIEF



Before THOMAS, KRASS, and BARRETT, Administrative Patent Judges.
BARRETT, Administrative Patent Judge.

DECISION ON APPEAL

This is a decision on appeal under 35 U.S.C. § 134(a) from the final rejection of claims 1-5, 12-16, and 23-27. Claims 6, 17, and 28 have been canceled. Claims 7-11, 18-22, and 29-60 have been objected to as depending from a rejected base claim.

We reverse.

¹ Application for patent filed February 11, 2000, entitled "Optimizing Database Queries Using Query Execution Plans Derived From Automatic Summary Table Determining Cost Based Queries," which is based on and claims priority under 35 U.S.C. § 119(e)(1) from U.S. Provisional Application 60/134,745, filed May 18, 1999.

Appeal No. 2004-1538
Application 09/502,820

BACKGROUND

The invention relates to the optimization of queries using automatic summary tables in a database management system.

Claim 1 is reproduced below.

1. A method of optimizing a query in a computer, the query being performed by the computer to retrieve data from a database stored on the computer, the method comprising the steps of:

(a) identifying one or more automatic summary tables (ASTs) that overlap the query by matching definitions of the ASTs with requirements of the query, wherein the requirements of the query satisfied by an AST are encapsulated as a set of properties;

(b) enumerating one or more alternative query execution plans (QEPs) for the query, including at least one QEP that represents one or more access paths of the identified ASTs;

(c) assigning a cost to each of the alternative QEPs;
and

(d) choosing a most efficient one of the alternative QEPs based upon the assigned costs.

THE REFERENCES

The examiner relies on the following references:

Osborn et al. (Osborn)	6,026,391	February 15, 2000 (filed October 31, 1997)
Agarwal et al. (Agarwal)	6,370,522	April 9, 2002 (filed March 18, 1999)

THE REJECTIONS

Claims 1-3, 12-14, and 23-25 stand rejected under 35 U.S.C. § 102(a, e) by Osborn.

Claims 4, 5, 15, 16, 26, and 27 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Osborn and Agarwal.

Appeal No. 2004-1538
Application 09/502,820

We refer to the final rejection (Paper No. 19) (pages referred to as "FR__") and the examiner's answer (Paper No. 26) (pages referred to as "(EA__") for a statement of the examiner's rejection, and to the brief (Paper No. 25) (pages referred to as "Br__") and reply brief (Paper No. 27) (pages referred to as "RBr__") for a statement of appellants' arguments thereagainst.

OPINION

The invention

It will be helpful to discuss the disclosed invention, in particular, some of the terminology that appears in the claims. An "automatic summary table (AST)" is a table resulting from executing a query where the definition of the summary table is based on a "full select" statement (page 2, lines 9-13; page 29, lines 19-21). An AST example is the rich-employee AST which is defined by a SQL statement for records of employees which make over \$100,000 (page 10; Table 1, page 15). In the prior art, the AST could not be used by an optimizer to improve performance (page 2, lines 6-8). A "query" means to interrogate a collection of data such as records in a database. A query has a set of "requirements" that usually include at least the table or tables is the data coming from; the selection criteria, which is the matching condition or filter; and, which columns or fields in the tables are to be displayed or printed in the result. An example of query requirements is shown in Table 3 (page 20).

Appeal No. 2004-1538
Application 09/502,820

The limitation of "identifying one or more automatic summary tables (ASTs) that overlap the query by matching definitions of the ASTs with requirements of the query" can be explained by example. Table 1 (page 15) shows the properties and matching information associated with the rich-employees AST, with one predicate of "SALARY > 100,000," and Query 1 has predicates "SALARY > 100,000" and "LOCATION = PITTSBURGH," so the rich-employees AST overlaps Query 1 as to "SALARY > 100,000" (page 24, lines 16-23). That is, the rich-employees AST was defined using "SALARY > 100,000" and the matching step determines whether the AST was derived in such a way that it can be used as a starting point to satisfy the query (page 2, lines 18-20).

The limitation that "the requirements of the query satisfied by an AST are encapsulated as a set of properties" is shown by the properties in Table 1 (page 15), where the AST "properties" resulting from the matching phase include relational properties such as the tables referenced, columns supplied, expressions computed, predicates applied, unique keys and functional dependencies in effect, aggregation performed, and so on (page 14, lines 6-8).

A "query execution plan (QEP)" specifies the sequence of database operations used to satisfy the query (page 17, line 7). Each QEP represents an "access path" to a table or an AST. For example, QEP1 could represent an "index access" to the employee

Appeal No. 2004-1538
Application 09/502,820

table, where the index is on the salary column and is used with the predicate "SALARY > 100,000" to directly access records of employees whose salaries exceed 100,000 (Fig. 4A; page 18, line 21, to page 19, line 6) and QEP2 could represent a "full table scan" of the employee table where each record is accessed and qualified against the predicates "SALARY > 100,000" and "LOCATION = PITTSBURGH." The QEPs can be assigned a cost and the most efficient QEP can be selected.

The rejection and arguments

Appellants argue that Osborn does not teach "identifying one or more automatic summary tables (ASTs) that overlap the query by matching definitions of the ASTs with requirements of the query, wherein the requirements of the query satisfied by an AST are encapsulated as a set of properties." The examiner finds that the automatic summary table (AST) corresponds to Osborn's summary table detailed in Fig. 3, element 68, and column 6, lines 51-64, and finds that the matching limitation is taught at column 6, lines 20-35 (FR6; EA5). The examiner states that summary tables are an integral part of Osborn's database and that summary tables are shown in Fig. 3 and that Osborn teaches transmitting an SQL statement to a cost optimizer module which considers access paths to the requested data and estimates the relative cost of each execution plan (EA10-11). Appellants argue that Osborn merely describes generating a set of potential execution plans based on

Appeal No. 2004-1538
Application 09/502,820

the available access paths, and estimating the cost of each potential execution plan based on the data distribution and storage characteristics for the respective tables, clusters, and indexes to be used (Br6; RBr4). It is argued that nothing in Osborn describes matching definitions of the ASTs with requirements of the query in order to identify the ASTs that overlap the query and encapsulating the requirements of the query satisfied by an AST as a set of properties (Br6; RBR4).

Appellants also argue that Osborn does not teach "enumerating one or more alternative query execution plans (QEPs) for the query, including at least one QEP that represents one or more access paths of the identified ASTs." The examiner relies on column 6, lines 14-30, where it is said that Osborn teaches transmitting a SQL statement to a cost optimizer module that is associated with the database element 34 in Fig. 2, where the cost optimizer element 42 considers access paths to the requested data from the database and derives a plan for executing the query that is the most efficient as described at column 6, lines 23-29 (FR6; EA5). The examiner states that Osborn is directed to estimating query response time for database inquiries based on access paths (EA11). Appellants argue that Osborn merely describes generating a set of potential execution plans based on the available access paths of the query, but not based on access paths of ASTs identified as overlapping the query (Br6; RBr4).

Appeal No. 2004-1538
Application 09/502,820

Analysis

We agree with appellants that Osborn does not expressly or impliedly teach limitations (a) and (b) of claim 1. The relevant part of Osborn states (col. 6, lines 16-41):

The cost optimizer 42 considers available access paths to the requested data from the database 32, and derives a plan for executing the query that is most efficient based on statistics maintained in a data dictionary associated with the respective table(s), along with their associated clusters and indexes, accessed by the SQL statement representing the query 40.

In particular, the cost optimizer 42 generates a set of potential execution plans (not shown) for executing the respective query SQL statement (40) based on the available access paths, and estimates the relative "cost" of each potential execution plan based on the data distribution and storage characteristics for the respective tables, clusters and indexes to be used....

The cost optimizer 42 compares the estimated costs of the potential execution plans and returns the smallest estimated cost 44, along with a result set 45 representing the selected execution plan for the input query, to a query performance prediction ("QPP") module 46

Osborn does not teach the limitation of "(a) identifying one or more automatic summary tables (ASTs) that overlap the query by matching definitions of the ASTs with requirements of the query, wherein the requirements of the query satisfied by an AST are encapsulated as a set of properties" in claim 1. The table in Fig. 3 of Osborn, relied upon by the examiner, is a query history table that is used in assigning a cost; it is not an AST and is not an AST used in determining a QEP. While there is some determination of a match between the table in Fig. 3 and a query

Appeal No. 2004-1538
Application 09/502,820


(col. 7, lines 10-12), this is not an overlap between an AST and a query, nor does it meet the limitation that "the requirements of the query satisfied by an AST are encapsulated as a set of properties." Osborn's statement that "the cost optimizer 42 generates a set of potential execution plans (not shown) for executing the respective query SQL statement (40) based on the available access paths" (col. 6, lines 23-26) teaches "enumerating one or more alternative query execution plans (QEPs) for the query," but does not teach "including at least one QEP that represents one or more access paths of the identified ASTs" because it does not teach identifying ASTs that overlap the query. Osborn does disclose "assigning a cost to each of the alternative QEPs" and "choosing a most efficient one of the alternative QEPs based upon the assigned costs," except that the QEPs do not "includ[e] at least one QEP that represents one or more access paths of the identified ASTs" as recited in limitation (b). For these reasons, we find that the subject matter of independent claims 1, 12, and 23 is not anticipated. The rejection of claims 1-3, 12-14, and 23-25 is reversed.

Agarwal does not cure the deficiencies of Osborn with respect to the independent claims. Accordingly, the obviousness rejection of claims 4, 5, 15, 16, 26, and 27 is reversed.

Appeal No. 2004-1538
Application 09/502,820

In summary, rejections of claims 1-5, 12-16, and 23-27 are reversed.

REVERSED


JAMES D. THOMAS
Administrative Patent Judge


ERROL A. KRASS
Administrative Patent Judge


LEE E. BARRETT
Administrative Patent Judge

BOARD OF PATENT
APPEALS
AND
INTERFERENCES

Appeal No. 2004-1538
Application 09/502,820

GATES & COOPER LLP
HOWARD HUGHES CENTER
6701 CENTER DRIVE WEST, SUITE 1050
LOS ANGELES, CA 90045